Intelligent Robots and Computer Vision XVII, SPIE Proc. 3522, Boston, MA, November 1-5, 1998

New planetary rovers for long range Mars science and sample return

P. S. Schenker, E. T. Baumgartner, R. A. Lindemann, H. Aghazarian, D. Q. Zhu, A. J. Ganino, L. F. Sword, M. S. Garrett, B. A. Kennedy, G. S. Hickey, A. S. Lai, L. H. Matthies; Jet Propulsion Lab.; B. D. Hoffman, Massachusetts Inst. Technology; T. E. Huntsberger, Univ. So. Carolina

Jet Propulsion Laboratory, California Institute of Technology
4800 Oak Grove Drive/MS 125-224
Pasadena, California 91109-8099
paul.s.schenker@jpl.nasa.gov

ABSTRACT

There are significant international efforts underway to place mobile robots ("rovers") on the surface of Mars. This follows on the recent successful NASA Mars Pathfinder flight of summer 1997. In that mission, the 11+ Kg *Sojourner* rover explored a small 50 meter locale about its lander over a several week period. Future planned science missions of the Mars Surveyor Program are more aggressive, seeking to autonomously survey planetary climate, life and resources over multiple kilometers and many months duration. These missions will also retrieve collected sample materials back to a Mars Ascent Vehicle (*MAV*) for more detailed analysis on Earth. In support of these future missions we are developing and field testing new rover technology concepts. We first overview the design and initial operations of **SRR-1** *(Sample Return Rover)*, a novel 10 kg-class four wheel, hybrid composite-metal vehicle for rapid (10-30 cm/sec) autonomous location, rendezvous, and retrieval of collected samples under integrated visual and beacon guidance. SRR is a light 88x55x36 (LWH) cm^3 vehicle collapsing to less than one third its deployed field volume, and carrying a powerful, visually-servoed all-composite manipulator. We then sketch development of the **FIDO rover** (*Field Integrated Design & Operations*), a new 50+ kg, six wheel, ~100x80x50 (LWH) cm^3, high mobility, multi-km range science vehicle which includes mast-mounted multi-spectral stereo, bore-sighted IR point spectrometer, robot arm with attached microscope, and body-mounted rock sampling corer. Currently in integration phase, FIDO rover will first be tested in September, 1998, "MarsYard (JPL)" operations, followed by CY99 full-scale terrestrial field simulations of a planned Mars '03 multi-kilometer roving mission (*Athena*-based science rover payload), demonstrating remote science selection, autonomous navigation, *in situ* sample analysis, and robotic sample collection functions.

Keywords: mobile robots, 3D sensing, sensor fusion, space manipulators, robotic sampling, Mars sample return, planetary science

1. INTRODUCTION

There are significant international efforts underway to place mobile robots ("rovers") on the surface of Mars [1]. This follows on the recent successful NASA Mars Pathfinder flight of summer 1997. In that mission, the 11+ Kg *Sojourner* rover explored a small locale about its lander over a several week period; rover control was largely programmed, performed within lander line-of-sight (~50 meters), requiring frequent earth operator command interactions. Future planned science missions of the Mars Surveyor Program are far more aggressive, the goal being to autonomously survey planetary climate, life and resources over multiple kilometers and many months duration, optimizing use of available mission time and climatic conditions. Such "long-range rovers" will enable wide-area imaging and selection of science sites, spectroscopically-based chemical analysis of surface terrain and outcroppings, collection of loose soil and extracted rock samples, and characterization of local meteorology and many other science functions. This surface data collection is fundamental to understanding the formative processes of Martian geology, mineralogy, weather; a possible history of life and extant examples; as well as suitability of the Martian environment to future human habitation.

The scope of such rover-based science is broad, as are its possible engineering implementations. Robotic functions need not in fact be limited to rovers, and can indeed work as integrated mission elements with rovers – e.g. the use of lander-based arms for sample collection (planned for late CY'98 flight on the NASA *Mars Volatiles and Climate Surveyor* mission), drilling platforms, observational/sampling balloons, and sub-surface penetrators and explorers (augurs, moles, etc.); reference [2] gives a brief perspective on technical activities in these areas. One can envision more advanced scenarios in which such multiple, networked robotic assets combine functions to cover large Mars surface areas, cooperating to increase science diversity, sustaining a long term, autonomous robotic presence accessible by many users through remote-linked earth Internet. Current mission models are less complex, based on use of one rover or sequential missions carrying same. As one example, a first rover devoted to long range mobile science might precede the landing of second rover dedicated to sample retrieval. Alternatively, with sufficient mission launch resources, orbital communications (relay to earth) and rover payload/engineering capability, a single vehicle/mission might conduct science exploration, on-board sample caching, and retrieval of samples to an earth return ascent spacecraft (MAV).

2. SAMPLE RETURN ROVER (SRR-1)

The Sample Return Rover (SRR-1), shown below in **Figure 1**, was conceived as an approach to the "sample cache grab" problem. The function of SRR-1 is to quickly, robustly, and autonomously as possible, go from a landed spacecraft, find a nearby sample cache (assumed to reside on a science rover that has completed its mission, and sits dormant), and retrieve the cache to a MAV containment at the lander. The operational horizon of the sample grab can in principle be small, given that techniques are currently under development to allow precision landing within as little as a hundred meters. Thus, SRR could possibly communicate remote via the lander link (versus a rover uplink to orbiter or direct-to-earth link) and be under visual observation/reference of lander stereo cameras. The operational model we have developed is broader. We assume that SRR may need to transit over the horizon,

Figure 1. SRR-1 in simulated operations (JPL arroyo), beginning cross terrain navigation to science rover (LSR-1 [5])

starts with approximate knowledge of the cache site (referenced to the lander surface coordinates), provides its own on-board visual sensing, and maintains an accurate state estimate of vehicle location. During its cross terrain transits, SRR-1 must in real time capture and analyze a visual terrain map, detect and avoid hazards; as it approaches the target cache site, SRR-1 must detect the presence of same, accurately localize it (as to relative distance and orientation), and take a suitable inbound heading. Finally, once this "terminal guidance" phase into the cache is completed, SRR-1 must, in a cooperative robotic workspace, visually maneuver its arm to acquire and transfer the cache (see an example of this operation in lower left inset, **Figure 1**.)

SRR-1 System Description

Mobility:	4 wheels - all actuated (DC brushed/Maxon RE025) Skid steering (on variable footprint/strut angles) Articulated rocker/shoulder assembly (active control) 20 cm dia. Volume Efficient Deployable Rigid Wheels Self-deployable strut hinges (1:3 volume stowage)
Mass:	7.2 kg (<5.0 baseline, 1.0 active rocker, 1.0+ HD wheels)
Volume:	44 x 55 x 22 cm^3 stowed 85 x 55 x 36 cm^3 deployed
Manipulation:	MicroArm-2: 3 degrees-of-freedom with actuated gripping end-effector, 0.7 meter total reach, cache acquisition function
Navigation Sensors:	Forward-looking stereo camera pair (120 degree FOV) for obstacle detection and sample cache localization Spot pushbroom (SPB) laser/CCD active terrain sensor Manipulator-mounted goal camera (20 degree FOV) for long-range detection of science rover Tri-axial accelerometer package for tilt sensing Single-axis gyroscope for heading determination
Computing Platform:	PC104 80486 processor running at 66 MHz PC104 peripherals: framegrabbers, A/D, motion controllers
Operating System:	VxWorks 5.3 (Tornado)

We summarize the SRR-1 system above; in subsequent sub-sections, we describe SRR-1 key technical features and development approaches used to implement them. Note that SRR-1 is, in its baseline fixed-shoulder configuration (non-posable rocker angle), a 5.0 kg system (less the ~2.0 kg *MicroArm-2* [9]). The shoulder actuation/drive electronics add about 1 kg, (as does also implementing the vehicle with a class of dual purpose rugged-ized wheels applicable to 40+ kg class science rover design).

2.1 SRR Mechanical Design

The Sample Retrieval Rover, whose mobility/structure/perception & control is further summarized in **Figure 2**, is shown therein upper left with struts and wheels in one of the many possible stowed positions (launch and cruise configuration) measures approximately 440 x 550 x 215 mm and weighs just over 7 kg. Deployment of the running gear occurs as the rover is lifted from its stowed position. Energy for the deployment of the struts is stored in torsional springs that surround the hinge pin, and the hinge is locked open by means of a spring loaded latch. The rover's *Volume Efficient Deployable Rigid Wheel*, cf. [4, 5] and references therein) opens by means of energy stored in extension springs that also serve as the locking mechanism, once deployed.

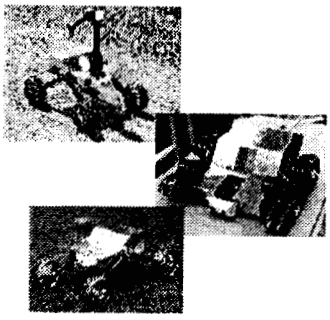

Mobility and Configuration Control:

- collapsible rocker with 3D composite latching torsional hinge
- enables stowage to 55x44x22 cm^3 from 55x85x36cm^3 (WLH)
- second generation VEDRW (.5 Kg/ea.): > 40 Kg load capability
- wheel actuation at 100:1 for 35 cm/sec, stall-torque of 14 N-M
- deployment mechanisms are self-energized with positive lock
- passive-bevel/active-spur-differential 180°-articulated suspension

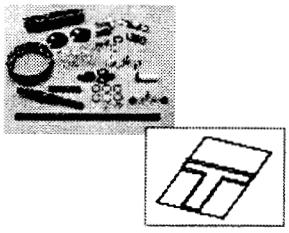

Materials-Structural-Thermal Design:

- T-sandwich panel/shear-plate paradigm shift from MPF/LSR-1
- graphite fiber/Al honeycomb in optimal load paths; > 16K in-lb
- compare: localized stiffness to minimized thermal conductivity
- Kevlar/graphite cyanate-ester VEDRW segments: 30% stiffer

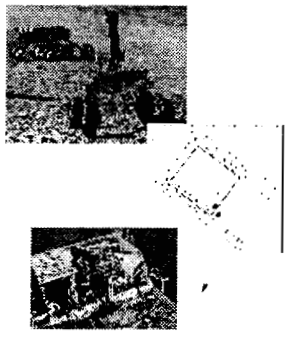

Perception & Autonomous Guidance:

- stereo (120° FOV) and goal acquisition (20° FOV) B&W imaging
- obstacle detection, visual terminal guidance, cache recognition
- VTG via localization of bar code skirt (2% range/.5 degree at 5 m)
- cache recognition via Wavelet-based "targeting" (1-3 mm at ~ 1 m)
- references cache cross-hairs to stereo workspace of MicroArm-2
- several stereo visual guidance & calibration paradigms developed

Figure 2. Summary of SRR-1 design features and performance parameters

The above noted deployable wheels, as also depicted in **Figure 3**, and wheel drive actuators have been designed to be compatible with a > 40 kg class rover (for dual use application to other JPL science rover designs). The Maxon RE025 DC brushed motor, combined with a custom designed 100:1 planetary gearbox, give SRR a top speed of 35 cm/sec, and stall torque of 14 N-M (resulting in "rim thrust" greater than the total weight of SRR in 1.0 gravity) with a life expectancy of many kilometers of actual testing (1 km of driving will result in only 160,000 revolutions of the motor). Some components of the wheels (spokes, hub, cleats, and "backbone") have also been redesigned (a net mass increase of 200+ grams per wheel as compared to a previous model used for LSR-1 [5]) so as to provide many kilometers of test life even on a vehicle significantly more massive than SRR.

Deployable hinges are fabricated from a new JPL-developed 3D composite material [3, 4], weigh 140 grams each, and enable the strut structure to be folded 180 degrees and passively "unstowed" by simply releasing restraints. Energy for deployment is obtained from torsion springs made from 2.16 mm diameter stainless steel wire wound "right-handed" 10 coils, with an outside diameter of 20 mm. Starting torque for deployment is approximately 1.4 N-M, with a final torque at full deployment of 0.5 N-M. The hinge is locked open when a spring loaded steel latch engages a steel dowel pin. There is no measurable "free play" in the hinge when it is in the locked position.

The articulated rocker assembly is a novel suspension system using a combination of a passive bevel gear differential bogie suspension and an actuated spur gear differential articulation joint. With a total system mass of approximately 0.8 kg, this suspension acts as a passive 4-wheel rocker suspension during normal driving while providing the added functionality of an articulated shoulder joint to aid various tasks such as stowage/deployment, active C.G. control, and configuration control. The passive suspension subsystem uses a composite main axle tube connected between each pair of struts and the bevel gear drive. The

chassis-mounted bevel gear housing supports the rover body and prevents excessive body rotation by distributing the angle of each pair of struts equally to the body, similar in nature to the passive suspension of Sojourner. A Maxon RE016 DC motor, attached to a custom 5-stage 1386.2:1 planetary gearhead and mounted internal to the axle tube on each side, powers a spur gear differential which controls the angle of each pair of struts. The spur gear differential is driven via a double-width pinion attached to the output of the gearhead. This differential consists of five main elements; in addition to the double-width pinion, there are the ring gears, housing, idler gears, and the doubler gears. The ring gears are 88 tooth, 64 pitch internal gears rigidly attached to each strut along the axis of the axle tube. The housing is axially mounted to, and fixed to rotate with the axle tube. Its main purpose is to maintain the position of the idler and doubler gears side-by-side on the double-width pinion. The idler gears are 34 tooth planets which interface the pinion to one ring gear and produce motion of the ring gear opposite in direction to that of the pinion. The doubler gears are two 17 tooth planets which interface the other half of the pinion to the ring and produce motion of the ring in the same direction as the pinion. The strut end fittings are constructed as a tongue and clevis joint so that when assembled the ring gears lie side-by-side. The idler, doubler and both sets of ring gears fit onto the double-width pinion and thus, make up the spur gear differential assembly. When the pinion rotates one strut will rotate in the same direction while the other in the reverse creating a scissors-like action between the struts. Moreover, this assembly compounds the gear ratio of the articulated system by another 4.4:1 bringing the total speed reduction (and consequently torque magnification) to approximately 6100:1. Factoring in assumed inefficiencies, this system delivers a maximum of approximately 32 N-m (282 in-lb) to each strut based on a stall torque of 0.0176 N-m (0.157 in-lb) for the RE016. With its tongue and clevis configuration, the articulated joint is capable of a full range of motion of 90° for each strut or a total of 180° included angle between the struts.

2.2 SRR Materials-Structural-Thermal Design

SRR includes several novel features in composite structural design. Previous rovers developed for the NASA Mars Pathfinder mission and under the Lightweight Survivable Rover task [5] have utilized a Warm Electronics Box [6] approach for the primary structural chassis. For those rovers, the WEB had to address both structural and thermal requirements. This led to a somewhat boxy implementation/encumbrance. In contrast, SRR's baseline design had structural requirements, but no thermal insulation requirements, nor the need of a solar array power source. This was predicated on early Mars Sample Return mission models for a *fast-cache-grab* (precision-land within 100's of meters, traverse the surface at highest possible speed under primary battery power, rendezvous and return the cache to the MAV for ascent within a single *sol* daylight period, thereby protecting MAV propellant from diurnal thermal cycling). SRR mechanical design is thus driven by need to provide a stable platform for the robotic arm and the mobility system loads. SRR changes the paradigm as to how these mechanical loads are distributed, by comparison to rovers

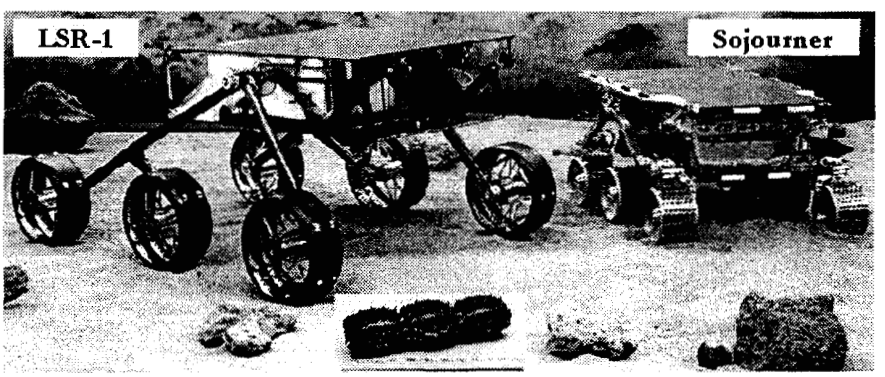

Figure 3. The JPL Lightweight Survivable Rover (LSR) is shown at left, a 7 kg vehicle having about 4X the operational volume of NASA's Mars Pathfinder flight rover, the 11 kg *Sojourner* vehicle shown to right. LSR, also a six-wheel rocker bogie mobility design, introduced a new *Volume Efficient Deployable Rigid Wheel* concept, illustrated lower center inset. This composite/Al VEDRW stows to ~ 30% its operational field volume and uses new 3D graphite-epxoy machined components.

previously developed at JPL. For the Sojourner rover, mechanical loads for the mobility system are all carried through the "jeff tube", and the WEB provides the mechanical support for the electronics and science instruments. As part of LSR-1 development the jeff tube was eliminated to significantly decrease its associated thermal losses, and the WEB was redesigned to take the mechanical loads for the mobility system and science instrumentation. For SRR there is a requirement that the robotic arm and the mobility system be rigidly integrated. To accomplish this, SRR utilizes a "Tee" sandwich panel/shear plate design. This design approach provides high localized stiffness for the load paths required for the mobility system shoulder joints and the robotic arm assembly. By using high stiffness graphite fiber facesheets and lightweight aluminum honeycomb sandwich structure selectively -- only in the load paths – we achieved a uniform 16,000 inch-lb stiffness. Contrast this with a localized 6,000 inch-lb stiffness in the Sojourner class WEB chassis design. In SRR, the specific stiffness was optimized compared to the specific thermal conductivity optimized in Sojourner.

Transferring all of the mechanical loads into a single plane thus removes the primary structural requirement from the rest of the SRR chassis. This allows the reduction in mass in the chassis. The rest of the chassis becomes an exoskeleton that is fitted to the volume of the electronics and interior mechanical assemblies. The SRR exterior shell is designed of lightweight E glass-epoxy with co-cured aluminized Kapton. The low emissivity aluminized Kapton provides a dual role as a thermal control surface and as the ground plane for the beacon assembly (as shown in **Figure 1**)

Another SRR materials innovation is use of hybrid Kevlar/graphite cyanate-ester laminates to fabricate the the collapsible wheel rim segments. This hybrid approach, using Kevlar as the outer surface, increases abrasion resistance of the wheel assembly and adds to wheel traction. The new wheel segments are 30% stiffer at the same mass than similar wheel segments in LSR graphite composite wheels. For SRR, we used 17-7 precipitation-hardened steel wheel cleats as were used on Sojourner (We are now completing a program of comparative wear/abrasion testing for a variety of simulated soil/rock surfaces and load conditions as a function of different cleat/material designs). The hybrid Kevlar/graphite concept was also used in wheel hub covers, as located on the inside of the wheel assemblies. These conical covers effectively prevent rock intrusion and fouling of wheel actuation/traction.

In general, SRR has made extensive use of the machinable 3D random fiber composite material developed in earlier work of our group [3, 4]. Applications include articulated shoulder joints, the deployable hinges, differential, camera bracket and other mechanical elements, including the robot arm (MicroArm-2 [9]) at large. This lightweight, high strength composite material affords up to a 30% mass reduction compared to equivalent 6061 aluminum designs. When this is combined with lightweight graphite composite structural tubes, the resultant mobility strut and robot arm assemblies possess both high strength and durability, and improved DTCE matching at critically stressed linkage interfaces

2.3 SRR Sensing and Visual Guidance

The overall sensing framework of SRR is cast in a state estimation paradigm based on Extended Kalman Filtering (EKF), as summarized in [7] and further reported by Baumgartner et al. in a companion paper of this symposium [8]. This development is complementary to other work we have recently reported on sensor fusion for robotic manipulation [9-11, and refs. therein]. In the broadest sense, our ***EKF fusion model*** embodies odometry, visual tracking of the rover's position/orientation, beacon heading-range, sun sensor data (celestially referenced surface position/inclination), gyro-heading, and other instrumentation (e.g., accelerometers). One illustrative example is shown at left, an overhead view of JPL's Planetary Robotics Laboratory test pit, with LSR-1 in progress to a goal under fused state guidance from odometry and incremental visual feature tracking [12]). Dead reckoning/odometry alone, can, as seen in this dramatic case, lead to large integral errors due to rover wheel slip and hang-ups. Similarly, loss of visual tracking (or limitations of calibration error in extended viewing fields can be problematic). Specific to visual guidance, there are several requirements we must address, as briefly recounted in the introduction to this section. These include 3D perception and terrain navigation, visual obstacle detection and avoidance, *visual terminal guidance* (visual goal acquisition and rendezvous of the rover in position/orientation with respect to the goal), and visual-based manipulation of the sample cache. SRR acquires 3D terrain estimates in two ways, one being a forward-looking ***Spot Pushbroom*** active sensor, which projects and reconstructs a lateral array if bright laser spots, and the other a conventional calibrated wide angle stereo pair (similar visual function to that which has been used in the JPL Rocky7 rover [13]). As we have previously reported of the Spot Pushbroom [5, 15], two small IR diode laser beams are projected into left-side/right-side 15 spot arrays via independent diffraction gratings and collimating optics, then imaged by and selectively read from a corresponding left-right CCD (monocular mode) camera pair. No difference imaging is required to detect the spots, even in bright sunlight (ROC ~.97). Based on the underlying camera/imaging model, spot locations are geometrically reconstructed at distances up to 1 m to accurately infer their 3D positions to ~ 1 cm cross-field resolution (in 512x486 high resolution readout mode). This enables real time coarse terrain analysis and hazard warnings, with the forward moving rover "sweeping" out significant terrain features and obstacle surfaces. Using this *coarse detection* mode, we safely attain 5-to-10 cm/sec *continuous* motion. The SPB acts as a front-end processor to trigger *fine detection* obstacle analysis by calibrated stereo [14], next described. SRR utilizes a *stereo camera subsystem* for dense range map generation. These range maps are used for reactive hazard avoidance navigation [15] as well as for localization the sample cache container relative to SRR's manipulator

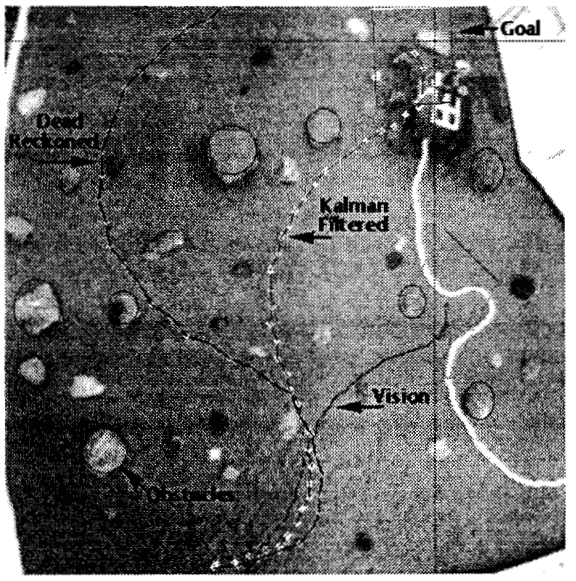

arm. The stereo camera pair consists of two B&W CCD-based imagers (Vendor: Supercircuts, P/N PC-8P) with 120 degree field-of-view lenses. The stereo baseline is 5 cm which results in a 2 cm range uncertainty at 1 m from the high-resolution stereo image pair of 512 x 486 pixels. For stereo-based hazard avoidance navigation, we have implemented an occupancy-grid technique for classifying the terrain directly in front of the rover. The range maps used to populate the occupancy grid are generated at a resolution of 96 x 60 pixels by down-sampling and smoothing a low-resolution stereo image pair of 256 x 243 pixels. Currently the occupancy grid covers an area which is 90 cm wide by 80 cm in front of the rover (the rover width is 55 cm). This area is broken up into 6-by-1 grid and the number of objects in each grid above or below a particular height threshold is computed. Based on this information, a look-up table is utilized to determine the steering angle that allows the rover to avoid the sensed obstacle field. The range map and steering angle computation takes approximately 4 sec on SRR's PC104-based 80486 processor. Therefore, we are able to realize continuous traverse speed of 5 cm/sec by ensuring that the first 20 cm of the occupancy map is always free of obstacles (and with further optimization, have the possible alternative fast navigation by stereo alone, although the SPB active sensing is attractive on other merits of low-light level operation, low feature terrain scans, etc.). We are in the process of implementing a denser occupancy grid encompassing a 100 cm by 100 cm area with a grid spacing of 5 cm by 5 cm. The rover will also keep track of the occupancy grids that were previously traversed to provide a local path planning capability.

One of the more challenging functions of the SRR operation is ***visual terminal guidance (VTG)***, wherein it must detect presence of a science target, establish a nominal heading and range to same, then carry out appropriate maneuvers and state updates to guide the rover into a terminal docking position (arm within a cooperative workspace). Our simplest implementation of this function, effective out to about 10 meters, is based on visual analysis and localization of 1-D cooperative features, as shown in **Figure 4.**

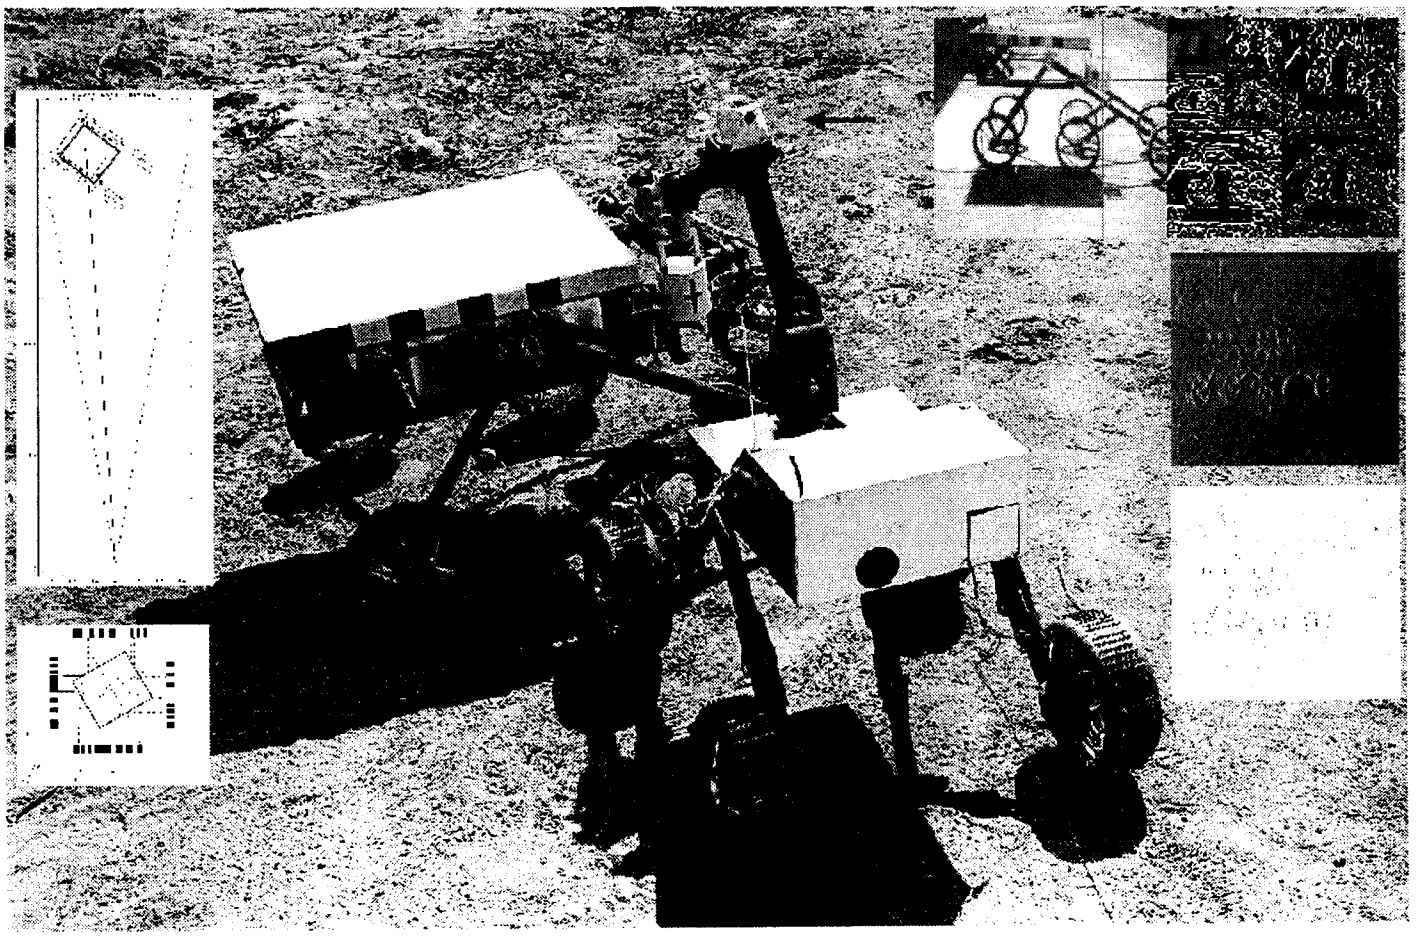

Figure 4. SRR docking with LSR for sample cache pickup (inset at left depicts geometric results of 1-D planar VTG analysis; results at right show intermediate image processing steps of cache detection (top), and VTG feature extraction)

In generality, SRR must detect a science rover in near proximity and arbitrary 3D perspective, estimating range-pose accurately in order to plan a traversal path to a sample cache container (possibly requiring mid-course corrections for obstacle avoidance, and goal reacquisition -- though retaining continuous visual track/servo on the target is far preferable). So as to minimize onboard computing complexity, we have first modeled a case wherein the two rovers (sample retrieval vehicle and cache site) are in a co-planar disposition (a reasonable operational assumption if the science rover termination site was selected, versus artifact). The general localization concept is one of having a reasonable fused state estimate of cache site distance and heading, activating a visual search with a specialized narrow-FOV goal camera, and basing the visual target detection process on use of cooperative site markings. Ideally, one might implement goal acquisition using inherent object features alone, but this is far more computationally intensive. We have designed a visual pattern – the quadrant-based bar code pattern of **Figure 4** lower left -- which allows us to reliably estimate range-pose parameters at 3-to-10 meter range with a 20° FOV monochrome camera (512x486 pixel resolution). Note that we in past have used a related scheme [16] to localize a rover with a colored cylinder; by comparison to techniques described here, that earlier method required a color camera (implicitly modeling the case of a stationary lander camera tracking a nearby rover) and implemented feature segmentation through a region-growing method, versus linear feature extraction analysis used here (wherein sub-pixel edge detection yields very good accuracy). Our visual pattern is attached as a "skirt" which fits around the solar panel edge of a rover (LSR being the target science rover). The visual pattern is designed to facilitate feature extraction. We implement an edge-based procedure. Each stripe is 1.5" (~3.75 cm) high and 2" (~5 cm) wide; a distinct edge that can be seen at a distance of 10 m. A wider stripe is placed at each rover corner so that edges from the background can be distinguished and edges from two sides can be separated. Taking advantage of the fact that both rovers lie on a flat terrain, only 1-D edge detection is applied, a derivative-of-Gaussian process. Then non-maxima suppression is applied and edge points are grouped according to the uniformity of edge strength. Finally, only edgels which satisfy geometrical constraints, i.e., similar length and equally spacing and polarity constraint, alternating positive and negative, are extracted as stripe candidates. To suppress noise, we only examine edges with magnitude in the top few percentile. After the stripe edgels are identified, sub-pixel edgels can be obtained through the parabolic interpolation.

We want to estimate the rover range-pose from one perspective-projected image and known goal camera and target models. Assume that we have matched corner points from the skirt model to the goal camera image. We then have an estimation problem, informally, $\mathbf{u_k} = \underline{\mathbf{P}} * [\underline{\mathbf{R}} * \mathbf{r_k} + \underline{\mathbf{T}}] + \mathbf{W}$, where $\mathbf{u_k}$ is understood to be the image coordinate pair projections of rectangular corner target patch locations $\mathbf{r_k}$, $\underline{\mathbf{P}}$ is the perspective projection, $\underline{\mathbf{R}}$ is the rotation transform, $\underline{\mathbf{T}}$ is the translation between imaging and target model coordinates, and $\mathbf{W}$ an unspecified noise process [14]. The geometric formulation is straightforward, and we assume an orthographic projection model. Worst performance is limited by the detectability of one set of stripes at extreme angles. Consider a worst case estimate for range of 4 m (the nominal SRR stand-off distance from science rover site prior to final approach/docking) where image resolution for our camera setup is 0.275 cm/pixel: For the stripes to be detectable, the stripe separation in the image should be at least be 2 pixels, inferring pose estimation error up to 6°, and range error of 10%. However, in the more favorable angular pose, errors are held to < 2° and <3% respectively (having in practice seen considerably better performance). More recently we have reformulated the problem for a case of generalized 3D central perspective, thus allowing us treat accurately cases of significant relative inclination between the two rovers, and accurate estimation at closer distances.

Using the VTG-based estimates above, we can accurately bring SRR into a cooperative workspace with respect to the science rover. This process requires a tangent-angle turn-away and vector onto the cache quadrant (cache is mounted in left rear of LSR-1 target vehicle) at about 4 m stand-off. Ideally, we maintain a visual-track and servo of SRR state during this phase of operations. E.g., during turn-away on the tangent circle, there is risk of losing goal acquisition, and badly degrading the SRR state estimate. This is particularly true if the only local position update is skid-steered odometry. This situation can be ameliorated by visual tracking of local features in the terrain (see photo example of page 6 above) for continuous guidance, as well as counter-steering of the arm-mounted goal camera to compensate SRR motion. These various data sources are input to the earlier mentioned EKF predictive estimator. This generalized capability becomes particularly important when terrain around the cache is pathological, requiring path planning and iterative course corrections to finalize the docking. Once SRR is in the cooperative workspace, we can exercise several options to locate the cache. The most basic is teleoperative, wherein SRR stereo views are downlinked for analysis and determination of the cache location. This may consume at least a sol, utilizing valuable mission time. Alternatively, cache localization can be autonomous, based on detection of features/markings in the stereo pair, establishing a stereo correspondence of same, solving for 3D cache position, and commanding appropriate inverse kinematics (IK) motion of the rover arm to modeled location of the cache pickup point. This eye-to-hand coordination problem [14, 17] is of course well-established in the robotics/computer vision literature and has been treated by us in several ways elsewhere [9, 10]; one potentially problematic aspect is quality of available stereo pair calibration and that of the related IK modeling, given the mechanical shock and thermal extremes these subsystems experience in flight and diurnal cycling. To this end, we have recently developed several approaches

[11, 18] to self-calibrating stereo and image plane based manipulation that require only approximate knowledge of the camera and arm parameters. As far as detecting the cache itself, we again use cooperative markings (the cross-hairs of **Figure 4**), and are exploring several different techniques of feature extraction, including a wavelet-based textural analysis [19], for which an example is shown upper right in **Figure 4**), and eigenface (composite basis function) models, which can in principle be made relatively robust to viewing aspect and background. We have achieved integrated localization/manipulation performance of 1-2 cm at 70 cm by conventional stereo IK positioning, wherein primary error appears attributable to camera calibration and predict better than 3 mm using our "calibration-adaptive" camera-space based methodologies (see related performance achieved in [9, 18]).

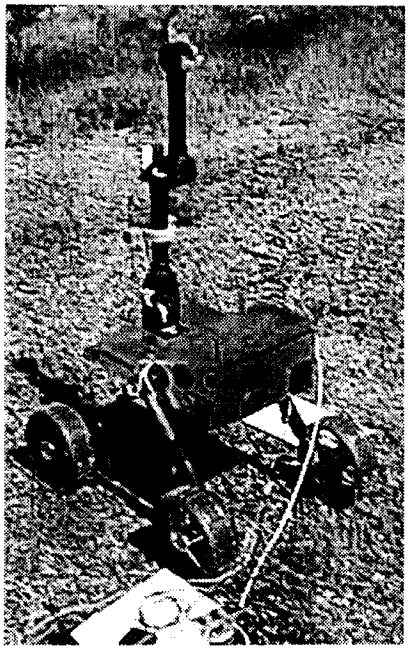

In closing this section we comment briefly on our development of beacon-based rover guidance. It is in principle possible to carry out the sample rendezvous functions (retrieval of the sample cache; and return to the MAV) under visual navigational guidance alone. E.g, given high-resolution landing descent imagery that is well-referenced to lander coordinates, and a feature based analysis of same, then SRR might visually track to-and-from a cache site. Or, suppose the sample return is direct from one science mission (the same rover performs science selection/cache return), and a single rover maintains and navigates by high quality map of its traversal from the lander; alternatively, if science activities are very local (10-to-100 m line-of-site about the lander), then lander imaging might close a visual servo loop to rover guidance. However, in all scenarios, particularly those beyond direct lander observation, and traversing complex paths to multiple sites, it is beneficial and expedient to have a cache beacon to provide a heading reference. The radio homing beacon that we are currently evaluating on SRR is of 'SuperDF' type, a proven design widely used by search and rescue teams and amateur radio enthusiasts. This device allows locating a transmitter of a selected frequency broadcasting CW, voice, or data. The phase detection design is resistant to the high noise environment of a rover, insensitive to signal strength, and provides an unambiguous direction indication. The direction finder uses two dipole antennas, a FM narrow band receiver, and some JPL custom electronics (The prototype unit uses a SuperDF unit purchased from BMG Engineering, Inc., Temple City, California, and a commercial FM 2 meter band receiver). The antennas are switched at 400 Hz, allowing the RF signal from one antenna only to pass to the FM receiver at any given time. When one antenna is closer to the transmitter, the signals arrive the two antennas at different times, resulting in is a phase difference between the signals. The FM receiver sees a sudden phase change at the instant that the antenna signals are switched, and interprets the phase change as a modulated signal. The receiver output is a brief pulse of polarity determined by the direction of the phase change, with magnitude proportional to the phase change. Analog switches allow only the phase pulse to pass while locking out other information which may contaminate the signal. Polarity and magnitude information is processed into a +/- 10V signal connected to the computer A/D port, enabling continuous target tracking during a traverse without the need to stop periodically to acquire a new heading. Current operating configuration allows for fully automatic operation by allowing the rover computer to switch the transmitter on when a new reading is required. Output linearity is a design-dependent variable. The response is naturally sinusoidal for an angle to the transmitter from -90 to +90 degrees. The system employed on SRR was adjusted to a linear region between -30 to +30 degrees, with greater angles indicating the direction to turn with an error of greater than 30 degrees. Recent field trials on open terrain (JPL's Arroyo, pictured in **Figure 1**) demonstrated the ability of the rover to accurately track the beacon signal from 100 meters to actual "collision" with the transmitter. Any tracking error was so small that the rover stayed within one meter of centerline along a narrow, one-lane dirt road from start to finish. Maximum operational distance has yet to be determined and is estimated to be greater than one kilometer, dependent primarily on transmitter power and the ability of the receiver to discriminate between signal and noise.

3. "FIDO" ROVER

As noted in the introduction, future Mars missions set a far more ambitious agenda for mobile science than that seen in the seminal Mars Pathfinder (MPF)/Sojourner flight. The goal there, exceeded in actual experimentation, was brief local lander operations under daily uplink/downlink of rover commands and science data. Sojourner traversed in line-of-sight about the MPF lander, navigated only by local hazard avoidance (coarse active sensing by laser striper) and dead-reckoning, carried a single rear-mounted instrument (the Alpha Photon X-ray Spectrometer) of long integration time, communicated via the lander, and baselined a week-long mission. By comparison, Mars rover missions of the '03/05-and-beyond time-frame are cast in a *"mobile science laboratory (MSL)"/in situ science* paradigm. Such planetary explorers will be capable of long range, over-the-horizon, self-navigating

traverses, communicating by direct-to-orbit (or even earth) relay. Onboard science instrumentation will be diverse, enabling visual identification of spectrally and morphologically interesting terrain areas, confirmatory analysis and target selection, microscopic visual and elemental/structural spectroscopic assays, and sample extraction and containment. Engineering instrumentation will include viewing masts, dexterous arms, and micro-coring devices. To this end, we have begun the development of an advanced technology vehicle intended for terrestrial prototyping and field simulation of Mars mobile science laboratory mission scenarios.

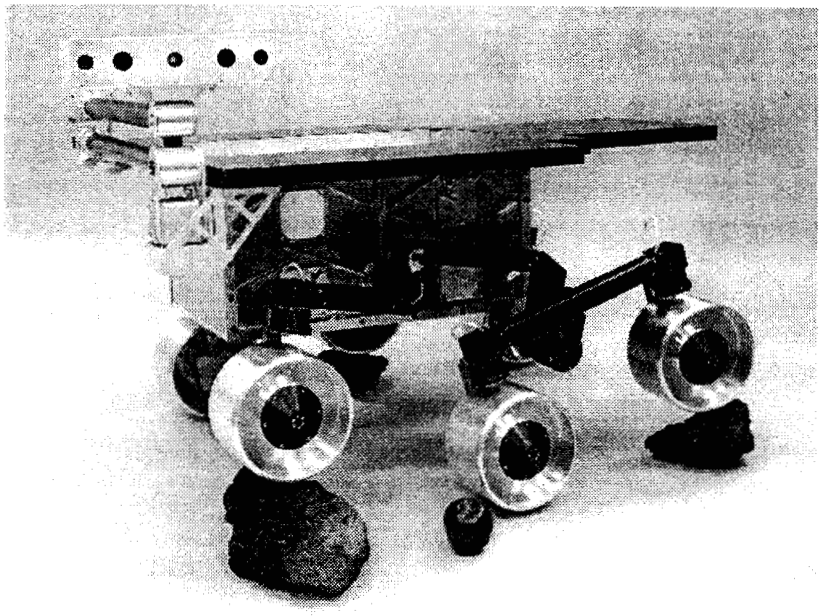

FIDO (*Field Integrated Design & Operations*) rover, shown at left in initial stages of mechanical integration, is a 50+ kg, six wheel, ~100x80x50 (LWH) cm^3, high mobility, multi-km range science vehicle with mast-mounted multi-spectral stereo, bore-sighted mid-IR point spectrometer, robot arm with attached microscope, and body-mounted rock sampling mini-corer. We list FIDO rover's initial configuration below. FIDO will first be tested in September, 1998, "MarsYard" local area operations at JPL, followed by CY99 full-scale desert field trial simulations of Mars '03 multi-kilometer roving mission functions (per the planned Athena science payload, cf. S. Squyres, Cornell University, PI). FIDO trials will be an experimental progression, remotely controlled by Web command/visualization interface [20]. Operation objectives include remote science selection, autonomous navigation, *in situ* sample analysis/manipulation, and robotic sample collection and containment functions.

Mobility Sub-System

—6 wheel rocker-bogie: each wheel independently driven and steered (35 N-m torque at stall); flight-related actuator design

—speed < 9 cm/sec, 20 cm dia. wheels, ground clearance of 23 cm

—rover dimensions of 1m (L) x 0.8m (W) x 0.5m (H); 68 kg mass

—4 degrees-of-freedom mast with integral science instrumentation

—4 d.o.f. instrumented science arm with actuated gripper/scoop

Navigation and Control

—PC104 platform, 80586 133 MHz processor, and I/O boards (frame-grabbers, digital I/O, serial I/O, A/D, motion control)

—front/rear hazard avoidance stereo camera pairs (120° FOV)

—inertial navigation system (INS) and CCD-based sun sensor

—differential GPS for ground-truth (in field trial applications)

Science Instrumentation

—mast-mounted *multi-spectral (650, 740, 855 nm) stereo* camera pair (10° FOV); also used for long-range navigation

—*mast-mounted point spectrometer* (1.25 - 2.50 microns); *arm mounted color imager, Raman and Mössbauer spectrometers*; body-mounted *Mini-Corer sampling & caching subsystem*

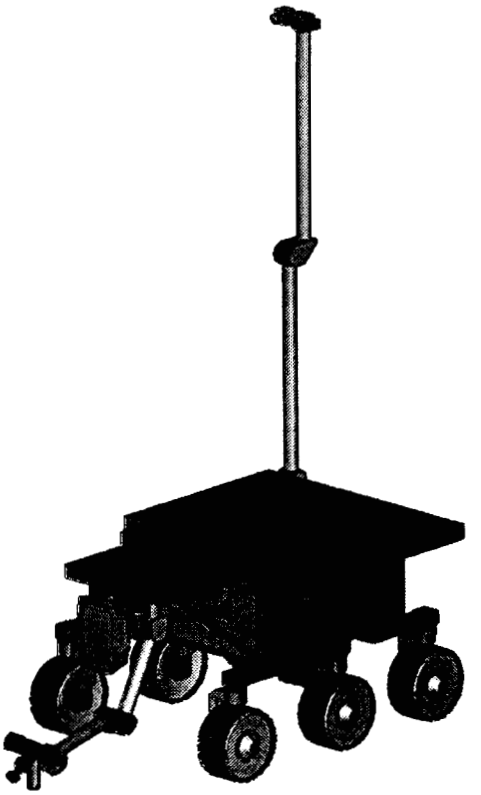

3.1 FIDO Computing and Software Environment

The primary computing hardware environment for the FIDO rover is based around the PC104 bus. This environment was chosen due to its low-cost, availability of multiple commercial I/O boards, and ease of upgrade. FIDO carries the following:

133 MHz 586 processor and operating system

- 32 Mbytes DRAM
- NE2000 Ethernet controller
- Four COM serial ports and bidirectional LPT port
- VGA CRT local bus controller
- 48 digital I/O lines
- 384 Mbyte IDE Flash Solid State Drive
- VxWorks 5.3 operating system

Baseline analog and digital I/O capabilities:

- Two black & white framegrabbers (512 x 486 and 256 x 243 image resolutions)
- One color framegrabber (640 x 480 image resolution)
- Two sixteen channel, differential-mode, 12-bit A/D converter board (0-5 V, 0-10 V, +/-5V input ranges)
- One eight channel, 12-bit D/A converter board (0-5 V, 0-10 V, and +/-5 V output ranges)
- Two 15 channel encoder reading boards (developed by Microcomputer Systems, Inc.)
- One D/A multiplexer board (developed by Microcomputer Systems, Inc.)

The control of all actuators is accomplished under microcomputer control via optical encoders feedback. All motor encoders are read through a custom-built frequency/phase measurement board with control signals sent to each motor amplifier as generated by the D/A converters. Each of the eight D/A outputs is multiplexed to the motors through a sample/hold circuit. The control loop runs at approximately 100 Hz.

Software development approach/drivers

Code is written in an object-oriented framework so that each code segment is easily added or removed from the software system. While specific to the hardware and instrument suite on FIDO, software must be capable of being easily modified to account for new hardware or instrument configurations. The code is resident on the rover's solid state disk and capable of working in a stand-alone mode (e.g., with VxWorks resident on the solid state disk, the rover code must be load-able on the rover--not require download from a remote host). In the stand-alone mode, the rover can be operated by a ground-control laptop PC running *WITS* [20]. No particular code segment requires more than 8 Mbytes of RAM, and all such core code segments are to be written in ANSI C, as are also all low-level software drivers. We chose the ANSI C specification to facilitate portability, compilation to a high-speed object code, minimal object code size, and high and low-level programming capability). Low level software drivers include B&W and color frame-grabbing, A/D, motor control (D/A, D/A mux, encoder reading), serial I/O, and digital I/O.

3.2 Navigation and Engineering Sensors

The FIDO rover carries three pairs of stereo navigation cameras: one forward-looking stereo pair, one rear-looking stereo pair, and one mast-mounted stereo pair. The forward and rear-looking stereo pair are wide field-of-view, B&W imagers (120x50 degrees horizontal/vertical) mounted at 50 cm height and are primarily used for reactive, obstacle avoidance maneuvers. The mast-mounted stereo pair (1.8 m height) is a color imaging system with a horizontal field-of-view of 10 degrees (IFOV ~ .28 mrad). The color imaging system is implemented by mounting a Liquid Crystal Tunable Filter (LCTF) as the front-end filter on a B&W imager. The LCTF is tuned to three frequencies in the near-IR spectrum (750, 850, and 950 nanometers). All rover actuators include optical quadrature encoders; steering actuators also include absolute potentiometers. Absolute potentiometers are also mounted to the rocker-bogie suspension system with one potentiometer on each bogie pivot and one potentiometer on the rocker pivot. The rover also carries a full inertial navigation system (INS). This includes three single-axis gyroscopes mounted in a tri-axial configuration and three single-axis accelerometers also mounted in a tri-axial arrangement. The inertial system will provide knowledge of the rover pitch, roll, yaw, distance traveled, speed, and acceleration. The accelerometers (P/N 7591-2) are +/- 2g sensors and are manufactured by Endevco, Inc. The gyroscopes (P/N QRS11) are quartz rate sensors and are manufactured by Systron Donner Inertial Division. We have also developed a CCD-based sun sensor to provide an independent estimate of rover heading.

FIDO Rover Functions

The following is the FIDO rover functionality that we are developing/integrating to carry out planned science operations of the FY98-end MarsYard engineering tests and subsequent CY'99 desert field trials at Silver Lake, CA (Mars'03 mission simulations):

Rover Navigation
- rover localization
  - wheel odometry
  - inertial navigation system
  - sun sensor
- stereo-based hazard avoidance
- way point navigation and path tracking

Health Monitoring and Fault Detection
- rover state data collection
  - accelerometers
  - gyroscopes
  - potentiometers
- power consumption
- motor current monitoring
- temperature monitoring

Instrument Arm Control
- instrument placement
  - forward and inverse kinematics
  - vision-based navigation
  - contact sensing
- science instrument interfaces and data collection
  - Color microimager
  - Mössbauer spectrometer
  - Raman spectrometer

Mini-Corer Control
- Mini-Corer placement
- Mini-Corer control and sample acquisition
- storage of sample in cache container

Mast Control
- mast positioning
  - forward and inverse kinematics
- acquisition of stereo panorama
- LCTF control for multi-spectral imaging
- point spectrometer
  - instrument pointing
  - instrument interface and data collection
- vehicle self-inspection

Communications
- parsing of commands to the on-board command processor
- telemetry outputs (images, science instrument data, engineering data)

Ground Operations
- Web Interface for Telescience [20]
- science data collection
- engineering data collection

ACKNOWLEDGEMENTS

This work was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. We gratefully acknowledge support of the NASA Telerobotics and Exploration Technology Programs for this research. Areas of cognizance for the above reported work are: P. Schenker, PI; A. Ganino, B. Kennedy, and L. Sword (now with IS Robotics, Inc.), mobility; G. Hickey, thermal controls & materials; D. Zhu and L. Matthies, sensing & hazard avoidance; E. Baumgartner, system engineering & integration; H. Aghazarian, software development & computing; M. Garrett, electronics; and A. Lai, science and data interfaces. B. Hoffman completed the visual tracking work noted herein as part of an in-residence MIT Master's Thesis study (P. Schenker NASA advisor, J-J. Slotine, MIT Faculty Advisor); T. Huntsberger carried out his work under a NASA Summer Faculty Fellow appointment (NASA/JPL host, P. Schenker).

REFERENCES

1. The Rover Team [D. L. Shirley at al.], "The Pathfinder Microrover," Journal of Geophysical Research, Vol. 102, No. E2, pp. 3989-4001, Feb. 25, 1997; see also http://www.jpl.nasa.gov [ref: Mars Pathfinder Mission, Sojourner Rover].

2. C. R. Weisbin, D. Lavery, and G. Rodriguez, "Robotics technology for planetary missions into the 21st century," Proc. 1997 Intl. Conf. on Mobile Planetary Robots, Santa Monica, CA, January 29-31 (ref: Planetary Society, Pasadena, CA, Org: L. Friedman), and, D. L. Shirley and J. R. Matijevic, "Mars rovers: past, present, and future," Proc. Princeton Space Studies Inst. 20th Anniversary Conf., May, 1997; see also, http://robotics.jpl.nasa.gov, and link therein to Rover and Telerobotics Program Overview (Mars technologies).

3. P. S. Schenker, Y. Bar-Cohen, D. K. Brown, R. A. Lindemann, M. S. Garrett, E. T. Baumgartner, S. Lee, S.-S. Lih and B. Joffe, "A composite manipulator utilizing rotary piezoelectric motors: new robotic technologies for Mars in-situ planetary science," in Enabling Technologies: Smart Structures and Integrated Systems, SPIE Proc. 3041, San Diego, CA, March, 1997.

4. L. F. Sword, G. D. Harvey, and J. A. Frazier, "Deployable rigid wheel," NASA Tech Briefs, June, 1997, ref: NPO-19982; D. K. Brown, "Nearly isotropic, resin transfer molded composites," NASA Tech Briefs, May, 1998, ref: NPO-19918.

5. P. S. Schenker, L. F. Sword, A. J. Ganino, D. B. Bickler, G. S. Hickey, D. K. Brown, E. T. Baumgartner, L. H. Matthies, B. H. Wilcox, T. Balch, H. Aghazarian and M. S. Garrett, "Lightweight rovers for Mars science exploration and sample return," Intelligent Robotics and Computer Vision XVI, SPIE Proc. 3208, 13 pp., Pittsburgh, PA, Oct. 14-17, 1997.

6. G. S. Hickey, D. Braun, L. C. Wen, and H. J. Eisen, "Integrated lightweight structure and thermal insulation for Mars rover," Proc. SAE 25th International Conference on Environmental Systems, San Diego, July, 1995; G. Hickey, D. Braun, L.C. Wen, and H. Eisen, "Integrated thermal control and qualification of the Mars rover," Proc. SAE 26th International Conference on Environmental Systems, Monterey, CA., July, 1996; G. Hickey, R. Manvi, and T. Knowles, "Phase change materials for advanced Mars thermal control," Proc. SAE 26th International Conference on Environmental Systems, Monterey, CA., July, 1996.

7. E. T. Baumgartner and S. B. Skaar, "An Autonomous Vision-Based Mobile Robot," IEEE Transactions on Automatic Control, Vol. 39, No. 3, pp. 493-502, March, 1994; J.-D. Yoder, E. T. Baumgartner, and S. B. Skaar, "Initial Results in the Development of a Guidance System for a Powered Wheelchair," IEEE Transactions on Rehabilitation Engineering, Vol. 4, No. 3, pp. 143-151, September, 1996; also, S. I. Rumeliotis and G. A. Bekey, "An extended Kalman filter for frequent local and infrequent global sensor data fusion," in Sensor Fusion and Decentralized Control in Autonomous Robotic Systems (Eds., P. S. Schenker, G. T. McKee), SPIE Proc. 3209, pp. 11-22, Pittsburgh, PA, Oct. 14-17, 1997.

8. E. T. Baumgartner, P. S. Schenker, Jet Propulsion Lab.; B. D. Hoffman, Mass. Inst. Technology; T. E. Huntsberger; Univ. So. Carolina, "Sensor fused navigation and manipulation from a planetary rover," in Sensor Fusion and Decentralized Control in Robotic Systems (Eds., P. S. Schenker, G. T. McKee), SPIE Proc. 3523, November, 1998, Boston, MA.

9. P. S. Schenker, E. T. Baumgartner, S. Lee, H. Aghazarian, M. S. Garrett, R. A. Lindemann, D. K. Brown, Y. Bar-Cohen, S. S. Lih, B. Joffe, and S. S. Kim, Jet Propulsion Laboratory; B. H. Hoffman, Massachusetts Institute of Technology; T. Huntsberger, Univ. of So. Carolina, "Dexterous robotic sampling for Mars in-situ science," Intelligent Robotics and Computer Vision XVI (Ed. D. Casasent et al.), SPIE Proc. 3208, 16 pp., Pittsburgh, PA, Oct. 14-17, 1997.

10. P. S. Schenker, D. L. Blaney, D. K. Brown, Y. Bar-Cohen, S-S. Lih, R. A. Lindemann, E. D. Paljug, J. T. Slostad, G. K. Tharp, C. E. Tucker, C. J. Voorhees, and C. Weisbin, Jet Propulsion Lab.; E. T. Baumgartner, Mich. Tech. Univ.; R. B. Singer, R. Reid, Univ. of Arizona, "Mars lander robotics and machine vision capabilities for *in situ* planetary science," in Intelligent Robots and Computer Vision XIV, SPIE Proc. 2588, Philadelphia, PA, October, 1995.

11. S. Lee, B. Hoffman, E. Baumgartner, and P. Schenker, "GOBS: An intelligent system architecture for planetary robotic sampling," in Proc. Intl. Conf. on Advanced Robotics (ICAR'97), Monterey, CA, July, 1997; see also, S. Lee, S. Ro, and P. Schenker, "Self-calibration of eye-hand coordination system with decentralized data fusion," in Sensor Fusion and Decentralized Control in Autonomous Robotic Systems, SPIE Proc. 3209, Pittsburgh, PA, October, 1997.

12. B. Hoffman, E. Baumgartner, P. Schenker, and T. Huntsberger, "Improved Rover State Estimation in Challenging Terrain", to appear in Autonomous Robots, February; Z. Zang "Iterative Point Matching for Registration of Free-Form Curves and Surfaces", International Journal of Computer Vision, 13:2 119-152 (1994); Z. Zang, "A Stereovision System for a Planetary Rover: Calibration, Correlation, Registration, and Fusion", Proc. IEEE Workshop on Planetary Rover Technology and Systems, Minneapolis, Minnesota, April 23, 1996; H. J. S. Feder, and J.-J E. Slotine, "Real-Time Path Planning Using Harmonic Potentials In Dynamic Environments," 1997 IEEE International Conference on Robotics and Automation, Albuquerque, NM, April.

13. S. Hayati, R. Volpe, P. Backes, J. Balaram, R. Welch, R. Ivlev, G. Tharp, S. Peters, T. Ohm, R. Petras, and S. Laubach, "The Rocky 7 rover: a Mars sciencecraft prototype," Proc. 1997 Intl. IEEE Conf. on Robotics and Automation, Albuquerque, NM, April 20-25; see also, R. Volpe, J. Balaram, T. Ohm, and R. Ivlev, "The Rocky 7 Mars rover prototype," Proc. IEEE/RSJ International Conf. on Intelligent Robots and Systems (IROS '96), Osaka, Japan, November 4-8, 1996.

14. R. Y. Tsai, "A Versatile Camera Calibration Technique for High-Accuracy 3D Machine Vision Metrology using Off-the-Shelf TV Cameras and Lens", IEEE Trans. on Robotics and Automation, Vol. 3, No. 4, pp. 323-344, 1987.

15. L. Matthies, T. Balch, and B. Wilcox, "Fast optical hazard detection for planetary rovers using multiple spot laser triangulation," Proc. 1997 Intl. IEEE Conf. on Robotics and Automation, Albuquerque, NM, April 20-25; L. Matthies et al., "Obstacle detection for unmanned ground vehicles: a progress report," Proc. 7th Intl. Symp. Robotics Research, Germany, October, 1996 (ref: Springer-Verlag, New York, NY, 1996, ISBN 3-540-76043-1); and, L. Matthies, C. Olson, G. Tharp, S. Laubach, "Visual localization methods for Mars rovers using lander, rover, and descent imagery," Intl. Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS '97), Tokyo, Japan, July, 1997.

16. L. Matthies, E. Gat, R. Harrison, B. Wilcox, R. Volpe, T. Litwin. ``Mars Microrover Navigation: Performance Evaluation and Enhancement", Autonomous Robots Journal, Special Issue on Autonomous Vehicle for Planetary Exploration, 2(4), 1995.

17. L. E. Weiss, A. C. Sanderson, and C. P. Neuman, "Dynamic Sensor-based Control of Robots with Visual Feedback", IEEE Trans. on Robotics and Automation, Vol. 3, No. 5, pp. 404-471, 1987; B. Espiau, F. Chaumette, P. Rives, "A New Approach to Visual Servoing in Robots," IEEE Trans. on Robotics and Automation, Vol. 8, No. 3, pp. 313-326, 1992; S. B. Skaar, W. H. Brockman, W. S. Jang, "Three-Dimensional Camera Space Manipulation", Intl. Jrnl. of Robotics Rsch., Vol. 9, No. 4, pp. 22-39, August, 1990.

18. E. T. Baumgartner and P. S. Schenker, "Autonomous Image-Plane Robot Control for Martian Lander Operations", Proceedings of the IEEE International Conference on Robotics and Automation, Vol. 1, pp. 726-731, Minneapolis, MN, April 1996; E. T. Baumgartner and N. A. Klymyshyn, "A sensitivity analysis for a remote vision-guided robot arm under imprecise supervisory control," in Sensor Fusion and Distributed Robotic Agents, Proc. SPIE 2905, Boston, MA, November, 1996.

19. T. Huntsberger, "Semi-autonomous multi-rover system for complex planetary surface retrieval operations," in Sensor Fusion and Decentralized Control in Autonomous Robotic Systems, SPIE Proc. 3209, Pittsburgh, PA, October, 1997; F. Espinal, T. Huntsberger, B. Jawerth, and T. Kubota, "Wavelet-based texture analysis for automatic target recognition," in SPIE Optical Engineering, January, 1998.

20. P. Backes, G. Tharp, and K. Tso, "The Web Interface for Telescience (WITS)." Proceedings of the IEEE International Conference on Robotics and Automation, Albuquerque NM, April 20-25, 1997.